

## View From the Boundary

BARBARA WEBB

*Center for Cognitive and Computational Neuroscience, Department of Psychology,  
University of Stirling, Stirling FK9 4LA, Scotland, UK*

**Abstract.** Re-implementing biological mechanisms on robots not only has technological application but can provide a unique perspective on the nature of sensory processing in animals. To make a robot work, we need to understand the function as part of an embodied, behaving system. I argue that this perspective suggests that the terms “representation” and “information processing” can be misleading when we seek to understand how neurobiological mechanisms carry out perceptual processes. This argument is presented here with reference to a robot model of cricket behavior, which has demonstrated competence comparable to that of the insect, but utilizes surprisingly simple central processing. Instead it depends on sensory interfaces that are well matched to the task, and on the link between environment, action, and perception.

### Introduction

The intersection of biology and robotics—the position of my own research—is often characterized as taking information from neuroethological investigations of natural systems to implement as new technology for man-made systems. However, another aspect of work in this area is to use the robotic implementations as a means of exploring biological hypotheses (Webb, 2000). This approach can provide a perspective on fundamental issues that is complementary to the view of the biologist engaged in primary research on the animal. This includes ideas on the most promising routes by

which biological understanding might inform technological developments.

My main thesis will be that examining invertebrate sensory systems from this perspective teaches us that they do not actually do much “information processing” or “representation”—depending, of course, on how you define these terms (see below). When we look at invertebrates, it appears that the function of the sensory systems is not to inform the animal generally but to control specific behaviors; that the means by which they do so is often determined as much by peripheral sensory physics as by central computation; and that appreciating the problem in terms of an embodied animal interacting with an environment is more appropriate than approaching it in terms of building an internal representation of the external stimuli. Wehner (1987) used the term “matched filters” to describe how animals may be faced with problems that apparently need sophisticated information processing solutions, but actually solve them by exploiting sensor mechanisms and behaviors that are uniquely matched to the required tasks. Further examples presented in this collection of symposium papers included the simple visual variables exploited by the bee to control flight (Srinivasan, 2001), and the use of “fanning” by crayfish (Breithaupt, 2001) or moth (Ishida, 2001) to improve chemical plume tracking.

Given that the terms “representation” and “information processing” are nevertheless commonly used by invertebrate neuroethologists (e.g., during the symposium Thomas Cronin discussed the scanning movements of the mantis shrimp eye as implying a relatively sophisticated system for registering the information properly onto a subjective representation of space [Cronin and Marshall, 2001]), a distinction may need to be drawn between this usage and the kind of full-blown symbolic encoding and manipulation that characterizes the “information processing” view of perception and cognition in traditional Artificial Intelligence. The claim that a pattern of neural firing represents a stimulus is

E-mail: b.h.webb@stir.ac.uk

This paper was originally presented at a workshop titled *Invertebrate Sensory Information Processing: Implications for Biologically Inspired Autonomous Systems*. The workshop, which was held at the J. Erik Jonsson Center for the National Academy of Sciences, Woods Hole, Massachusetts, from 15–17 April 2000, was sponsored by the Center for Advanced Studies in the Space Life Sciences at the Marine Biological Laboratory, and funded by the National Aeronautics and Space Administration under Cooperative Agreement NCC 2-896.

often only a claim that the firing and the presence of the stimuli causally co-vary: in the same way the electric current in a wire might be said to represent the position of an on-off switch. However, some deplore this as a misuse of the term "representation." For example, Maze (1981; p. 87) says that "the connection between the brain state and the external fact the knowledge of which it subserves . . . is just that of cause and effect, not representation," and Clancey (1991; p. 110) argues that "structures in the brain that cannot be perceived [by the agent] have no representational status to the agent."


The background for this disagreement over appropriate usage reflects two distinguishable senses of the relationship of representation. The first I will call "intentional representation," defined in the theory of signs by Peirce (cited in Fetzer, 1988) as "something that stands in for something (else) in some respect or other for somebody" (p. 134)—for example, use of the term "LGN" by a scientist to represent a part of the brain. The second I will call "causal representation," which describes indirect or mediated presentation, for example, the activity of ganglion cells presenting retinal stimulation patterns to LGN. The critical distinction between these two is that the "intentional" case requires that the thing represented can be directly experienced by the representor: the scientist can hear the sound "LGN," or look at the brain part, and this is why he can use one to represent the other. In the "causal" case the LGN cannot access the retinal activity independently—for example, to confirm that the "representation" by the ganglion cells is correct. To illustrate the distinction another way: an ant may use a pattern of landmarks as a representation of a nest position, in which case it can know about the presence of the landmarks and the presence of the nest in the same way (*i.e.*, through its senses). If the ant is also said to use the response pattern of neurons in its brain to "represent" the presence of the landmarks, the ant's relationship to the neural firing and to the landmarks are not comparable. We are using different levels of description when we say it "recognizes" the landmarks or "recognizes" the pattern of neural firing. (A possible source of confusion here is that looking at the ant's behavior and its neural processes from our point of view, we may well find that one (the firing) seems to stand in for the other (the landmark): but this is "intentional representation" only to the experimenter; to the ant it is merely "causal representation.")

Similarly, there are distinctions to be drawn between usages of "information processing." There is the formal communication theory sense as defined by Shannon (1948); there is the everyday sense in which information is taken to be something containing meaning; and then there is the more recent identification of information processing with computation—that is, involving syntactic manipulation. None of these maps directly onto the usage whereby, for example, lateral inhibition in the retina is called information

processing (there is no well-defined sender, receiver, or probability function; the meaning is opaque in the same way that the "representation" is non-intentional; and the processing is governed by physical rather than syntactic rules). The more apposite term here would seem to be *signal* processing, but "information processing" has become ubiquitous.

Do these distinctions matter, or are they mere semantics? I would argue they are important because the explanatory power of applying the terms is very different. It is an empirical, and somewhat controversial, hypothesis to say that invertebrate behavior is controlled by intentional internal representations, manipulated in meaningful information processing. Whereas to say that behavior is controlled by "causal" representations and involves "signal" processing is merely to say that the activity of the nervous system has a role in controlling the behavior, which was not in doubt. The same point has been expressed by Beer (2000, p. 97) with regard to cognitive science: "If any internal state is a representation and any systematic process is a computation then a computational theory of mind loses its force."

Moreover, it is not always clear that insect neuroethologists, in their usage, are not drawing conclusions that rest on conflating the meanings. An example is the tendency to start from the observation that an animal behaves differently in the presence of some stimulus, go on to describe the process involved as the animal internally "identifying" that stimulus before responding to it, and from this end up looking inside the brain for the neural mechanism that carries out the "identification." If the use of "identify" is only metaphorical, then it should not constrain the interpretation of findings, but it does. As an illustration, we can consider a classic piece of neuroethology in cricket phonotaxis research, the discovery of "recognition" neurons in the cricket brain whose firing rate response corresponds remarkably well to the likelihood of tracking by the cricket when it is presented with songs of different syllable rates (Schildberger, 1984a). Although this discovery is certainly of significance in trying to disentangle the neural wiring underlying the behavior, this "representation" of the "attractiveness" of syllable rates by the firing rate of an identified neuron is by no means an explanation of the behavior. First, it is not surprising, given that the animal behaves in different ways to different songs, that we find some neurons active under conditions when it does respond and not active when it does not—this is simply to say that its motor behavior is under some kind of neural control. Furthermore, the result does not in itself tell us how the neuron comes to have this property: understanding the mechanism of "recognition" requires understanding the neural connectivity leading to this property, which to date is still not fully resolved. Finally, in the neural model described below, we found highly comparable property of correlation of firing rate in certain neurons with syllable rate preference—yet the firing rate here had no functional role in the behavior but was simply a side-effect (Fig. 1). In fact,



**Figure 1.** The “firing rate” of a neuron in the robot model matches the “phonotactic preference” displayed in behavior. This looks like the “recognition” neuron discovered in the cricket (Schildberger, 1984) but in fact plays no functional role in the behavior. (Adapted from Webb and Scutt, 2000.)

there need not be any explicit “identifier” in the brain for the animal to single out and approach a specific signal, as I will now describe in more detail.

### Modeling Cricket Behavior

Cricket phonotaxis—the ability of females to track down male calling songs—includes a significant range of the problems of responding appropriately to specific sensory signals: identifying the signal against a noisy background; recognizing that it is the correct one; localizing its source; possibly choosing between rival signals. An information processing approach to this problem identifies the problems to be solved by the cricket’s neural system as filtering for the right carrier frequency and filtering for the right repetition rate to recognize the signal (Popov and Shuvalov, 1977; Thorson *et al.*, 1982; Stout and McGhee, 1988); comparing the amplitude of the auditory signal between two sensors to determine the direction of the source or at least which way to turn (Schmitz *et al.*, 1982; Schildberger and Horner, 1988; Huber, 1992); and separating simultaneous sound sources sufficiently to assess and approach the more attractive one (Doherty, 1985; Simmons, 1988; Pollack, 1998).


However, closer examination of the peripheral sensing system in the animal suggests that it may solve at least some of these problems directly, without any explicit representation of the song. The pressure difference receiver mechanism that enables the animal to detect the sound direction (Michelsen *et al.*, 1994) is inherently dependent on that sound being within a particular range of wavelengths. The neural encoding of the subsequent intensity difference between the ears is potentially in the form of a temporal code

(Schildberger, 1984b; Stumpner *et al.*, 1995) that could explain the pattern dependency of the response. Finally, the animal’s behavior in response to sound will position it in the sound field in such a way that it is likely to end up at the most attractive source rather than confused between them (see below). In other words, the behavior does not require any internal representation of the nature or position of the sound source.

That this is indeed possible has been demonstrated in a robot implementation of this suggested mechanism for phonotaxis (Webb, 1995; Lund *et al.*, 1997; Webb and Scutt, 2000). The robot has an auditory system that, like the cricket’s ears, uses cross-delay and summation of the two signals to produce a strongly directional response despite small receptor separation. Because the delay is fixed, the wavelength of the signal is a crucial determinant of the effectiveness of the device. Thus the robot will, for example, locate a 4.7-kHz signal better than one at higher or lower frequencies, and will preferentially approach a 4.7-kHz signal when a song of differing frequency is simultaneously presented, with no other form of frequency filtering.

The behavior of the robot is controlled by a spiking neural network consisting of only four units. Two input units integrate the auditory signal and initiate firing above a threshold (their behavior is closely modeled on the response properties of identified neurons [AN1] in the cricket). They respectively excite two output units, but cross inhibit each other’s axons. Thus the unit that fires first effectively suppresses the effect of the other side. The input-output connection is further modulated by synaptic suppression—that is to say, successive spikes have progressively less effect on the postsynaptic membrane potential, unless there is a gap in which the synapse can recover. The result is that unless the input has an appropriate on-off pattern, it is not effective in generating an appropriate motor response as controlled by the output units. For example, the robot will show consistent tracking behavior only to songs that fall within a particular band of syllable repetition rates, the same as that preferred by the cricket (Fig. 2). Although this behavioral preference has a corresponding neural “representation” in the firing rates of the output units (Fig. 1), the actual explanation of the behavior lies in the interactions of the neural time courses of summation and decay, and indeed these generate the appropriate response much faster than the time that would be needed to get a reasonable estimate of the firing rate.

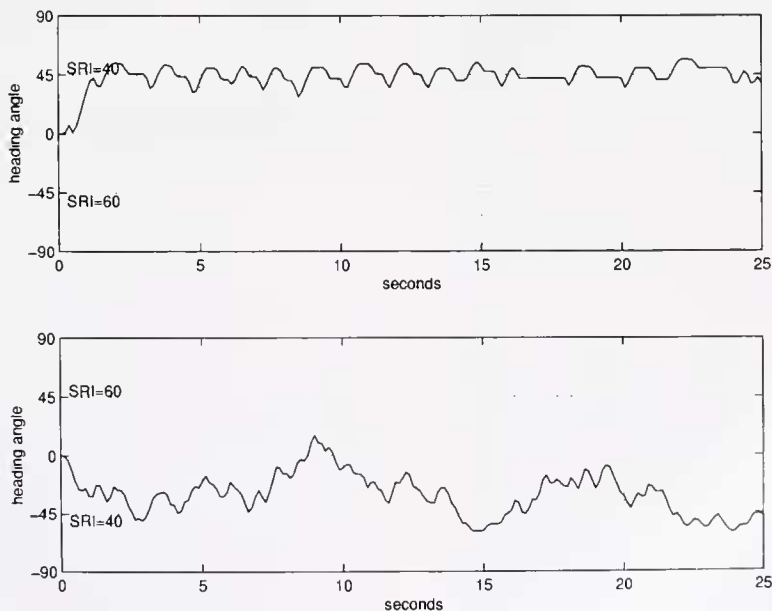
Having the model implemented in a physical device allowed us to test the behavior in realistic sound fields that would be difficult to simulate convincingly. Further characteristics of cricket behavior could thus be shown to emerge from the interaction of the controller, the physical interface, and the environment, without requiring further elaboration of the model. With sound from directly above (*i.e.*, lacking any horizontal directional difference), the



**Figure 2.** Tracking behavior of the robot in response to cricket songs at different syllable repetition intervals (SRI). The sound is at 45 degrees to the starting position of the robot. A SRI between 26 and 58 ms (comparable to the cricket) is needed for the robot to consistently turn and meander in the sound direction. (From Webb and Scutt, 2000.)

robot, like the cricket (Weber *et al.*, 1981), showed a tendency to perform tracking-like behavior without actually following one consistent direction. When the sound from above was paired with a continuous (*i.e.*, unattractive) stimulus from one side, the robot, like the cricket (Stabel *et al.*,

1989), tracked away from the lateral stimulus. When two similar sounds were played simultaneously, the robot could choose and track one of them because once it had turned slightly more to one side; the sound from that side captured the response. If the sounds differed slightly in temporal



**Figure 3.** Tracking behavior of the robot to simultaneous cricket songs at different syllable rates. The robot (like the cricket) turns and tracks the faster repetition rate (SRI = 40) whether it is on the left (upper plot) or the right (lower plot). (Adapted from Webb and Scutt, 2000.)

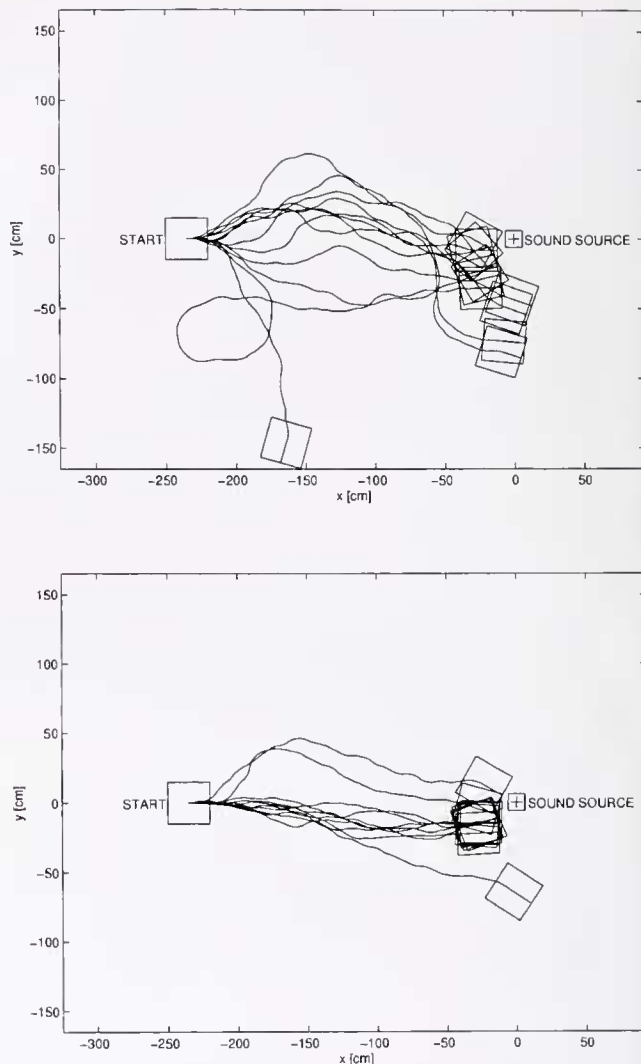
pattern, the robot, like the cricket (Doherty, 1985), could consistently choose one as the more attractive signal (Fig. 3).

### Integrating Sensory Systems

One argument advanced in favor of (real) information processing solutions is that they are more amenable to scaling up to explain more complex, flexible behaviors such as the integration of different sensory sources to control behavior. From an engineering or designer point of view, this might indeed be the case. Whether it is true of biology is another question: perhaps biological systems can offer us alternative schemes—perhaps more specialized to the animal's task niche, but on the other hand flexible and robust—for solving these kinds of problems. As a preliminary starting point for investigating these issues, I will describe some recent work done in collaboration with Reid Harrison (Webb and Harrison, 2000a,b) to look at the integration of the phonotaxis behavior on the robot with another fundamental sensorimotor reflex, the optomotor response.

Like many other insects, crickets will rotate in response to rotation of their visual surroundings. Normally this serves as a basic stabilization mechanism. The underlying sensor and neural circuitry for this response has been closely studied, particularly in the fly (Gotz, 1975; Reichardt and Poggio, 1976; Heisenberg and Wolf, 1988; Egelhaaf and Borst, 1993). It has been suggested that, in lit conditions, crickets will additively integrate their phonotaxis response and their optomotor response (Bohm *et al.*, 1991), which could improve the accuracy of their approach to sound (Weber *et al.*, 1981) by controlling for unintended course deviations.

A sensor that embodies the hypothesized mechanism of the optomotor response has been built in analog VLSI (very large scale integration) hardware (Harrison and Koch, 1998). This is a single chip that contains photoreceptors, temporal filters, comparison units, and widefield summation. The output can be used as a "torque" signal for the direction and approximate velocity of motion that would compensate for the visual rotation. We interfaced this chip to a robot that also had the sound-sensing circuit and neural model for phonotaxis described above. The two behaviors were initially combined in a directly additive way; that is, the motor output was a weighted sum of the signal given by the phonotactic turning decision and the signal given by the optomotor torque. However, this caused some problems, because turns in response to sound would generate strong visual rotation signals that the robot would attempt to correct, thus negating the initial turn. As a second approach, we used an inhibition scheme in which the robot would ignore the optomotor signal while turning in response to sound (other possible solutions are discussed in Webb and Harrison, 2000b).



**Figure 4.** Tracks of a robot with a 20% motor bias. Top, using phonotaxis only. Bottom, with an optomotor response added. The optomotor behavior significantly improves the ability to go directly to the sound source. (From Webb and Harrison, 2000b.)

With this simple interaction scheme it was possible to show that the added optomotor capability could significantly improve phonotaxis, more obviously so under conditions where the motor capability was made less reliable. Thus Figure 4 shows the behavior of the robot when approaching a sound source with an induced bias in its motor output that makes the left wheel turn 20% faster than the right. Without the optomotor response the robot had some difficulty reaching the speaker; with the response added it successfully and directly reached the speaker on all but one trial. Because the two hardware sensor systems are well tuned to executing their specific tasks, it was relatively simple to combine the behaviors to produce a robust performance without any explicit representation of the "fused" auditory and visual information.

## Conclusion

Robotics engineers already know a lot about information processing on representations. It is the standard computational paradigm, but it has proved difficult to employ to get robots to display behavioral competence comparable to even "mere" invertebrates. What they can learn from biology is how to build smart sensors that are matched to tasks; how to devise control systems that include patterns of behavior as part of the sensing process; and how to design internal nervous systems that exploit these factors. Calling these latter kinds of processes "representation" and "information processing" obscures the distinctive character of the mechanisms on offer. There is much yet to learn about the interplay of environments, behaviors, physics, and physiology. Biologists may have as much to learn from attempts to implement these mechanisms as do engineers.

## Literature Cited

- Beer, R. D. 2000. Dynamical approaches to cognitive science. *Trends Cognit. Sci.* 4: 91–98.
- Bohm, H., K. Schildberger, and F. Huber. 1991. Visual and acoustic course control in the cricket *Gryllus bimaculatus*. *J. Exp. Biol.* 159: 235–248.
- Breithaupt, T. 2001. Fan organs of crayfish: devices supporting chemical information exchange with the environment. *Biol. Bull.* 200: 150–154.
- Clancey, W. J. 1991. Situated cognition: stepping out of representational flatland. *Ai Commun.* 4: 109–112.
- Cronin, T. W., and J. Marshall. 2001. Parallel processing and image analysis in the eyes of mantis shrimps. *Biol. Bull.* 200: 177–183.
- Doherty, J. A. 1985. Phonotaxis in the cricket, *Gryllus bimaculatus* De Geer: comparisons of choice and no choice paradigms. *J. Comp. Physiol. A* 157: 279–289.
- Egelhaaf, M., and A. Borst. 1993. A look into the cockpit of the fly—visual orientation, algorithms, and identified neurons. *J. Neurosci.* 13: 4563–4574.
- Fetzer, J. H. 1988. Signs and minds: an introduction to the theory of semiotic systems. Pp. 133–161 in *Aspects of Artificial Intelligence*, J. H. Fetzer, ed., Kluwer, Berlin.
- Gotz, K. G. 1975. The optomotor equilibrium of the *Drosophila* navigation system. *J. Comp. Physiol. A* 99: 187–210.
- Harrison, R. R., and C. Koch. 1998. An analog VLSI model of the fly elementary motion detector. Pp. 880–886 in *Advances in Neural Information Processing Systems 10*, M. I. Jordan, M. J. Kearns, and S. A. Solla, eds., MIT Press, Cambridge, MA.
- Heisenberg, M., and R. Wolf. 1988. Reafferent control of optomotor yaw torque in *Drosophila melanogaster*. *J. Comp. Physiol. A* 163: 373–388.
- Huber, F. 1992. Behavior and neurobiology of acoustically oriented insects. *Naturwissenschaften* 79: 393–406.
- Ishida, H., T. Nakamoto, T. Moriizumi, T. Kikas, and J. Janata. 2001. Plume-tracking robots: a new application of chemical sensors. *Biol. Bull.* 200: 222–226.
- Lund, H. H., B. Webb, and J. Hallam. 1997. A robot attracted to the cricket species *Gryllus bimaculatus*. Pp. 246–255 in *Fourth European Conference on Artificial Life*, P. Husbands and I. Harvey, eds. MIT Press, Cambridge, MA.
- Maze, J. 1981. *The Meaning of Behaviour*. Allen and Unwin, London.
- Michelsen, A., A. V. Popov, and B. Lewis. 1994. Physics of directional hearing in the cricket *Gryllus bimaculatus*. *J. Comp. Physiol. A* 175: 153–164.
- Pollack, G. S. 1998. Neural processing of acoustic signals. Pp. 139–196 in *Comparative Hearing: Insects*, R. R. Hoy, A. N. Popper, and R. R. Fay, eds. Springer, Berlin.
- Popov, A. V., and V. F. Shuvalov. 1977. Phonotactic behaviour of crickets. *J. Comp. Physiol. A* 119: 111–126.
- Reichardt, W., and T. Poggio. 1976. Visual control in the fly. A quantitative analysis. *Q. Rev. Biophys.* 9: 311–375.
- Schildberger, K. 1984a. Multimodal interneurons in the cricket brain: properties of identified extrinsic mushroom body cells. *J. Comp. Physiol. A* 154: 71–79.
- Schildberger, K. 1984b. Temporal selectivity of identified auditory interneurons in the cricket brain. *J. Comp. Physiol. A* 155: 171–185.
- Schildberger, K., and M. Horner. 1988. The function of auditory neurons in cricket phonotaxis. I. Influence of hyperpolarization of identified neurons on sound localization. *J. Comp. Physiol. A* 163: 621–631.
- Schmitz, B., H. Scharstein, and G. Wendler. 1982. Phonotaxis in *Gryllus campestris* L. (Orthoptera, Gryllidae). I. Mechanism of acoustic orientation in intact female crickets. *J. Comp. Physiol. A* 148: 431–444.
- Shannon, C. E. 1948. A mathematical theory of communication. *Bell System Tech. J.* 27: 379–423.
- Simmons, L. W. 1988. The calling song of the field cricket *Gryllus bimaculatus* (de Geer): constraints on transmission and its role in intermale competition and female choice. *Anim. Behav.* 36: 380–394.
- Srinivasan, M. V., S. Zhang, and J. S. Chahl. 2001. Landing strategies in honeybees and possible applications to autonomous airborne vehicles. *Biol. Bull.* 200: 216–221.
- Stabel, J., G. Wendler, and H. Scharstein. 1989. Cricket phonotaxis: localization depends on recognition of the calling song pattern. *J. Comp. Physiol. A* 165: 165–177.
- Stout, J. F., and R. McGhee. 1988. Attractiveness of the male *Acheta domestica* calling song to females. II. The relative importance of syllable period, intensity, and chirp rate. *J. Comp. Physiol. A* 164: 277–287.
- Stumpner, A., G. Atkins, and J. F. Stout. 1995. Processing of unilateral and bilateral auditory inputs by the ON1 and L1 interneurons of the cricket *Acheta domestica* and comparison to other cricket species. *J. Comp. Physiol. A* 177: 379–388.
- Thorson, J., T. Weber, and F. Huber. 1982. Auditory behaviour of the cricket. II. Simplicity of calling-song recognition in *Gryllus* and anomalous phonotaxis at abnormal carrier frequencies. *J. Comp. Physiol. A* 146: 361–378.
- Webb, B. 1995. Using robots to model animals: a cricket test. *Robot. Auton. Syst.* 16: 117–134.
- Webb, B. 2000. What does robotics offer animal behaviour? *Anim. Behav.* 60: 545–558.
- Webb, B., and R. R. Harrison. 2000a. Eyes and ears: combining sensory motor systems modelled on insect physiology. Pp. 3913–3918 in *IEEE International Conference on Robotics and Automation*, San Francisco, 24–28 April 2000.
- Webb, B., and R. R. Harrison. 2000b. Integrating sensorimotor systems in a robot model of cricket behavior. *Sensor Fusion and Decentralised Control in Robotic Systems III*. SPIE, Boston, 6–8 Nov. 2000.
- Webb, B., and T. Scutt. 2000. A simple latency dependent spiking neuron model of cricket phonotaxis. *Biol. Cybern.* 82: 247–269.
- Weber, T., J. Thorson, and F. Huber. 1981. Auditory behaviour of the cricket. I. Dynamics of compensated walking and discrimination paradigms on the Kramer treadmill. *J. Comp. Physiol. A* 141: 215–232.
- Wehner, R. 1987. Matched filters—neural models of the external world. *J. Comp. Physiol. A* 161: 511–531.